c_o_l_e_c_c_i_ó_n
c_a_n_o_r_a III

M_I_N_I_N_G T_H_E C_O_M_P_U_T_A_T_I_O_N_A_L U_N_I_V_E_R_S_E

F_O_R M_U_S_I_C_A_L I_N_S_P_I_R_A_T_I_O_N

t_h_e f_o_u_n_d_a_t_i_o_n_s o_f m_y m_u_s_i_c_a_l l_a_n_g_u_a_g_e

Gustavo Díaz-Jerez

Ateneo de La Laguna

Also in this collection:

T_H_E S_O_U_N_D O_F A_N I_S_L_A_N_D

Canora collection. Original idea & coordination: Samuel Aguilar

English translation: Dácil Sánchez
Design: Samuel Aguilar, David Ángel Pérez Lage
Layout: David Ángel Pérez Lage
Editors: Orbelinda Bermúdez, Samuel Aguilar
Canora III: mining the computational universe for musical inspiration
Legal deposit: TF 2-2023
ISBN: 978-84-127680-1-5

M_I_N_I_N_G T_H_E C_O_M_P_U_T_A_T_I_O_N_A_L U_N_I_V_E_R_S_E

F_O_R M_U_S_I_C_A_L I_N_S_P_I_R_A_T_I_O_N

t_h_e f_o_u_n_d_a_t_i_o_n_s o_f m_y m_u_s_i_c_a_l l_a_n_g_u_a_g_e

Ateneo de La Laguna

INDEX

PROLOGUE

Samuel Aguilar

I had heard about Gustavo Díaz-Jerez and his achievements as a pianist and composer for years, actually decades. Of course, I was familiar with some of his work, but it wasn't until relatively recently that our paths finally crossed. However, perhaps due to everything I already knew (and I suppose he had heard something about me as well), at that moment, I had the feeling that even though we had never shared a conversation before, there were already certain previous connections that would make our mutual understanding smooth and natural. And indeed that was the case. The topics that attract Díaz-Jerez (which are seen not only in his discourse but also in his work) led to interesting conversations and shared perspectives with great ease.

There are composers who dedicate their lives to perpetuate and endlessly repeat the knowledge and traditions inherited from the past. On the other hand, there are others (perhaps more dangerous) who seek novelty without a foundation, without study, without the slightest interest in what has been done up to that point, in the paths that others have already travelled before them. It's novelty for novelty's sake. Díaz-Jerez is not at either of those two points. His musical work is based on rigour, knowledge, exquisite analytical skills, brilliant mathematical thinking, and a profound respect for tradition, while also demonstrating a strong interest in innovation. The eminent scientist Albert Einstein is believed to have made a statement in which he encouraged learning from the past, living in the present, and trusting the future. This could very well be a description of Gustavo Díaz-

Jerez's music. As you read through this notebook, you will see examples of works that pay tribute to composers from the past and that are even composed using their musical materials. However, you will also see how Gustavo uses the tools of the present and looks into the future by promoting and advocating the use of AI in composition.

We find ourselves in the Anthropocene, an era of absolute dominance by *Homo sapiens*. However, the advancements in AI are shaking many foundations. In recent years, fear has grown, and numerous opinions predict a dystopian future in which our species will be surpassed by machines created by us. Nevertheless, in line with what Díaz-Jerez argues in this notebook, renowned chess player Garry Kasparov recently stated in an interview, "The relationship between humans and machines has not been fully understood. Machines can inspire creativity because they offer new opportunities to reach our potential. (...) So instead of trying to compete with them, we should find the right algorithm to cooperate. That's why I strongly advocate collaboration with machines. I believe the best translation of AI is 'Augmented Intelligence.' It sounds friendlier. It doesn't conjure up those dreadful images of Terminator or The Matrix." It's clear that Gustavo Díaz-Jerez is indeed understanding this necessary relationship between machines and humans. In this notebook, although he poses some challenging questions, by sharing the foundations of his musical language, he exemplifies and demonstrates how AI can be effectively applied in musical creation.

INTRODUCTION

We all have an essence, a gaze fixed upon something specific that we have been drawn to since childhood. In my case, it has always been that which inherently possesses a structure, whether in nature or in various fields of science. Over time, this observation evolved into a rational understanding that I wanted to apply from my earliest steps as a composer.

I am fascinated by the process of unveiling mathematical structures that, through different means, I transform into musical material. This creative journey is akin to extracting a raw diamond from the depths of the earth and then carefully shaping and polishing it until it becomes a jewel. This endeavour is a blend of exploration in nature or within mathematical structures and a craftsmanship process, bridging the gap between scientific methods and artistic, musical goals, in my case.

My motivation lies in delving deep, first in pursuing, and then in meticulously modifying and shaping it, experimenting until it acquires a musical personality that I'm satisfied by. I don't aim to transform that raw diamond into a jewel that pleases an external audience. Of course, as the artist I am, I fell grateful for the positive response to the resulting composition, but it is not my primary objective in itself.

Throughout history, Western music has largely been based on four paradigms. Firstly, the diatonic model was predominant until the 16th Century. Gradually, this model gave rise to a second paradigm, the tonal system, which evolved by increasing chromaticism until its "collapse" in the early 20th Century. It was followed by free atonality, where the rules of tonal harmony were no longer upheld, and both sonic relationships and formal structures were created anew in each piece. The fourth paradigm, Schönberg's twelve-tone system, was seen as a logical step in the evolution of the tonal system, attempting to place the twelve semitones of the chromatic scale on a hierarchical scale of equality.

These paradigms provide timeless structures, serving as scaffolds that composers manipulate until the moment of their temporal inscription, that is, when a composition is created. These timeless structures include tetrachords, Greek modes, scales, rhythmic values, dynamic values, harmonic rules, as well as predetermined formal schemes like the sonata form. Composers have used them time and time again. They are, so to speak, the raw materials of music. These timeless structures provide a means through which composers can create variation and development.

I believe that aesthetic and value judgments are not applicable to these timeless structures, as they are mere abstract constructs from which real works of art are generated. They can be likened to the materials used by a sculptor (marble, wood, iron, etc.), the pigments used by a painter, or the words used by a poet.

The development of computer technology in recent years has provided composers with a way to incorporate new timeless structures into the compositional process. These structures stem from a discipline as ancient as humanity itself: mathematics. Mathematics and music have been intimately connected since the time of Pythagoras (circa 500 BC). Throughout the history of music, many composers have used mathematical models as a source of compositional creativity. However, only recently and thanks to technological advancements, composers can incorporate complex mathematical models into their works without having to perform the tedious and monotonous calculations they once required.

The objective of this article is to demonstrate that mathematical models provide composers with timeless structures that yield variation and development.

MAPPING

The connection between mathematical models and the timeless musical structures they generate (the musical raw material) is achieved through a process called mapping. Mapping involves establishing a correspondence (bijection) between the data of the mathematical model (usually numerical) and an ordered set (or sets) of timeless musical structures (scales, rhythmic values, dynamics, etc.). Mapping thus creates a direct link between the mathematical model and the set of musical parameters chosen by the composer.

There are essentially two types of mapping techniques: modulo-based mapping and normalised mapping.

Through module operation

A module operation is a straightforward mathematical procedure performed between two integers. It involves dividing the first number by the second and taking the remainder. For example: 10 MOD 3 = 1, because the remainder of dividing 10 by 3 is 1. Put more generally:

$$x \bmod y = n$$

When applying this mapping method to a numerical set, ***x*** represents the numerical value to be mapped, and ***y*** represents the total number of elements in our set of musical parameters. The result, ***n***, is therefore a number between 0 and ***y-1***, corresponding to the index of the element in our set of musical parameters. Below, I present an example that illustrates this type of mapping. We start with two numerical sets (vectors) that we are going to map, one for pitch, which we shall call **V**:

V: {0 0 2 0 5 4 0 0 2 0 7 5 0 0 12 9 5 4 2 10 10 9 5 7 5}

Where 0 corresponds to middle C (***C4***), 1 to ***C#4***, and so on up to 12, which corresponds to ***C5***. It's a vector of 25 elements with 13 different values (ranging from 0 to 12). The vector for the rhythmic values, which we'll call **R**, would be:

R: {2 0 3 3 3 7 2 0 3 3 3 7 2 0 3 3 3 3 3 2 0 3 3 3 7}

Where 0 corresponds to a sixteenth note, 1 to an eighth note, 2 to a dotted eighth note, and so on, up to 7, which corresponds to a whole note. This vector also consists of 25 elements with 8 possible values (ranging from 0 to 7). Next, we will apply the modulo operation to these two vectors:

x (the value to be mapped in the **V** vector) MOD 13 for the pitch set.
x (the value to be mapped in the **R** vector) MOD 8 for the rhythmic values set.

Once these operations are applied to each of the values in each vector, the resulting musical notation is as follows:

As we can see, it's the melody of 'Happy Birthday.' In other words, any musical structure can be encoded by a finite number of numerical vectors. Similarly, any numerical structure from the realms of mathematics or computing can be converted, through mapping, into musical structures.

Normalised

For processes in which the data set is limited to very small or fractional intervals, modulo mapping is not effective because it only considers the integer part of each value in the vector. In such cases, the most appropriate type of mapping is normalised mapping.

This mapping method uses the following formula to establish the correspondence (bijection) between the numerical data in the vector and the ordered elements of the musical set:

$$\text{Event No.} = \left[\frac{(\text{value} - \text{minval}) * \text{numevents}}{\text{maxval} - \text{minval}} \right]$$

where "value" is the current value of the numeric vector we want to map, "minval" is the minimum value in the vector, "maxval" is the maximum value in the vector, and "numevents" is the total number of elements in our musical parameter set. The brackets ([]) indicate that only the integer part of the result should be taken, disregarding the decimals. The "minval" and "maxval" values of the numeric vector must be known in advance. The result (***Event No.***) will be a number between 0 and "numevents-1," corresponding to the elements of our musical parameter set. Below, I propose the following example to illustrate this type of mapping. For simplicity, we will use a single musical set, limited to pitch, but we could extend it to rhythm, dynamics, etc.

The numeric vector to be mapped, which we'll call V, will be:

V: {0.892500, 0.342519, 0.803963, 0.562655, 0.878486, 0.381093, 0.842024, 0.474880, 0.890247, 0.348814}

The values of this vector are derived from the "logistic equation," a recurrence relation proposed in 1976 by the biologist Robert May:

$$x(n+1) = x(n)*\mu*(1-x(n))$$

For a value of μ (mu) equal to 3.57, the values of this equation are bounded within the interval [0, 1].

The set of musical parameters, which we'll refer to as **M**, consists of 24 elements, ordered from 0 to 23, and simplified for this case to represent the pitch of the sound, where 0 corresponds to ***C4*** (middle C), 1 to ***C#4***, and so forth up to 23 (which is ***G#5***). If we apply the normalised mapping to each element (***value***) of the previously described vector **V** using the formula for normalised mapping.

$$\text{Event No.} = \left[\frac{(\text{value} - \text{minval}) * \text{numevents}}{\text{maxval} - \text{minval}} \right]$$

Where "minval" is **0.348814** (the smallest value among the elements of **V**), and "maxval" is **0.892500** (the largest value), and "numevents" is 24 (the total number of elements in our musical parameter set), we obtain this new vector **Q** where each value represents the index of the corresponding element in the musical parameter set **M** described earlier:

Q: {21, 8, 19, 13, 21, 9, 20, 11, 21, 8}

Here is its musical transcription:

In this case, for simplicity, we have arbitrarily assigned eighth notes to the rhythmic dimension, but there's nothing preventing us from assigning the vector V to a set of rhythmic values, dynamics, and so on. Furthermore, it's possible to assign different numerical vectors from different mathematical processes to different musical dimensions and combine both types of mapping. The only limit is the composer's imagination. It's worth noting that performing all these calculations by hand would be incredibly tedious and prone to errors. Therefore, having some programming knowledge is necessary for these processes to be executed quickly and automatically by a computer.

We should never forget that these musical fragments, although they can become very extensive and complex, remain timeless musical structures, pure musical raw material that the composer must then work on, shape, and insert into the timeline of the composition within a musical narrative and coherence.

All of us composers are aware that embarking on the creation of a new piece always involves a certain level of stress and uncertainty, a sense of exploration, and to some extent, delving into the unknown. I would like to briefly describe how my creative process unfolds when approaching the composition of a new piece.

There are two elements that are usually predetermined, whether it's a commissioned work or something I initiate on my own. These two elements are the approximate duration and the instrumentation. It's important to note that these two elements substantially shape our creative process, as they affect aspects as relevant as form, timbral complexity, difficulty (writing for an orchestra is quite different from composing for a soloist), and more.

I always start with an initial inspiring idea. This idea can stem from a specific source, a literary work, mythology, and, of course, from purely scientific or mathematical concepts. This initial process involves "mining" the computational universe for structures that, through the aforementioned mapping process, provide me with musical material I am satisfied with. It's a highly exciting process because it always entails venturing into uncharted territory.

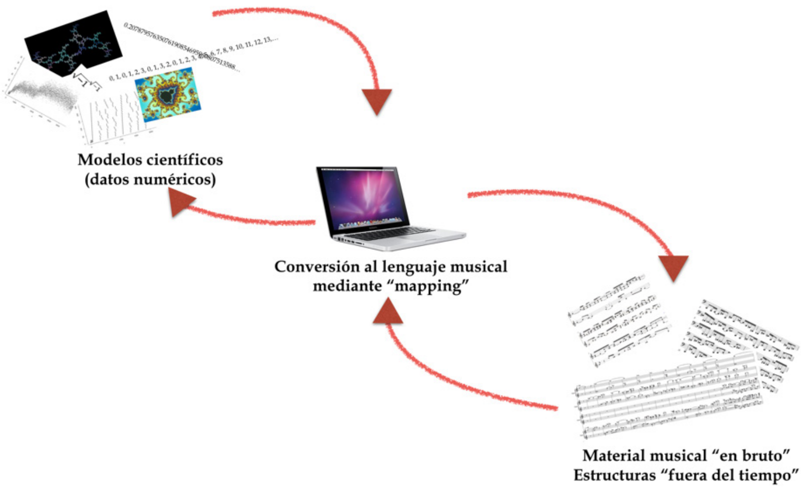

Once I find musical raw material with which I am satisfied, another equally thrilling process begins: adapting these timeless structures for performance, if the work is intended to be performed. Often, the generated material is extremely complex, and it requires a "pruning" process to simplify it, to make it "discrete" without losing its essence, and to adapt it to the technical possibilities of each instrument. In the case of purely acousmatic music, this adaptation process is not necessary. As a performer myself, I consider this adaptation process to be fundamental, as it largely determines whether the work will be embraced by performers and whether it will stand the test of time.

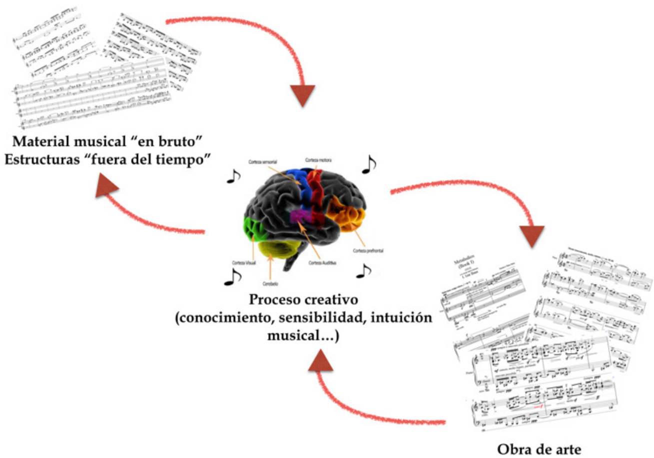

Finally, once this process of adapting the material is complete, the insertion of it into the timeline of the work begins. For me, this point represents the "starting gun." It's a process that is entirely "intuitive" and "manual" in the sense that, at least in my case, it no longer involves algorithmic or computational processes. It's a path that unfolds gradually, irreducible and unpredictable, influenced by a multitude of factors and presenting an infinite number of branching paths.

SUPPORT TOOLS FOR CREATION

During the search for musical raw materials, I use many tools, most of which are computer programmes that make this work possible. Among them, the most common ones are:

- ***FractMus 2000***

FractMus is a generative music program of my own creation, which I developed during my doctoral studies at the Manhattan School of Music in New York. It's a freely distributed (freeware) program for *Windows*. *FractMus* uses twelve algorithms from various branches of mathematics (number theory, cellular automata, chaos theory, etc.) that, through the mapping process, convert data into musical material. It contains numerous composition assistance tools. The program generates MIDI files that can then be loaded into editing software such as *Sibelius* or Finale for further refinement and processing.

Here are some works whose musical raw materials come from FractMus:

Kenotaphion (Metaludios for piano, Book I)

Hommage à H. Radulescu (Metaludios for piano, Book II)

Poetic and Useless Tool, for Ensemble.

- ***Apophysis***

Apophysis is a fractal image generation program (flame fractals). It is also freely distributed (freeware). I convert the images it generates into musical material through a process we will study later.

- ***Adiopaint***

Audiopaint is a freely distributed program for *Windows* that converts images into sound files (WAV). It offers multiple options that allow for precise control over the conversion process, including synthesis using sinusoids or granular synthesis, control over frequency, duration, and more. With this program, I generate audio from fractal images, which I can then convert into musical material.

- ***Melodyne***

Melodyne is a comprehensive commercial program with numerous tools for converting audio (WAV) to MIDI, allowing you to work directly with musical material in the score editor (*Sibelius*). Melodyne is a comprehensive commercial program with numerous tools for converting audio (WAV) to MIDI, allowing you to work directly with musical material in the score editor (*Sibelius*).

- ***Paul's Extreme Sound Stretch***

This powerful tool functions like a kind of slow-motion camera for sound. It allows you to stretch any sound while preserving its full timbral richness. The program is quite comprehensive, offering multiple conversion algorithms, filters, and more. It's freely distributed for *Windows*. I frequently use this program in combination with *Melodyne* to create musical raw material and to compose purely acousmatic music.

- ***Wolfram Mathematica***

Mathematica is a commercial program with built-in libraries for various areas of technical computing, enabling machine learning, statistics, symbolic computation, data manipulation, and more. It's also a powerful programming environment. I use it for manipulating various mathematical functions that I later incorporate into my compositions.

- ***Oeis.org***

The Online Encyclopaedia of Integer Sequences is a website created by Neil Sloane that contains hundreds of thousands of integer sequences organised and cross-referenced. It's a valuable resource for me in searching for musical material.

- ***Magenta***

Magenta is an open-source research project by *Google* that explores the role of machine learning in creative processes. Artificial intelligence is becoming increasingly integrated into our lives and is a field that can enhance human creativity. I use this platform to create AI models trained with my own music.

In addition to the mentioned programmes, I occasionally use other programmes like *Spear* (sinusoidal partial editing and resynthesis), *Sonic Visualiser* (audio visualisation and analysis), *AudioSculpt* (audio visualisation, analysis, and processing), *Max* (visual programming language for music and multimedia), *Pd* (a free alternative to *Max*), *Orchidea* (computer-assisted creative orchestration), *Open Music* (visual programming language for computer-assisted composition), etc.

Now, I will break down the main techniques I have developed over time to generate timeless structures that I later incorporate into my creative process

SELF-SIMILAR MOTIFS

Fractal structures

A fractal is a geometric object whose basic structure, often fragmented and irregular, repeats at different scales (self-similarity). Many natural structures (plants, coastlines, clouds, animal organs, neural networks, and even the large-scale structure of the universe) exhibit fractal characteristics. Fractal geometry is a branch of mathematics that formally describes these natural phenomena much better than traditional geometry. A typical mathematical fractal is the Sierpinski triangle (exact self-similarity), and a typical natural fractal is the Romanesco broccoli (approximate or statistical self-similarity).

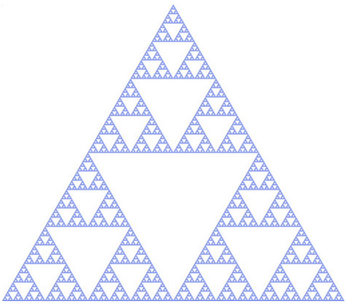

An important and consistent feature in my compositional language is the use of self-similar motifs. I'd like to demonstrate how I construct these motifs. This highlights the exploration, the mining within the computational universe that I mentioned earlier. First, I begin with a numerical vector, which can be a mathematical constant, astronomical data, and so on, for example:

$$\mathrm{Sinh}(e) = 7.5441371028169758\ldots$$

With Wolfram Mathematica, I convert this constant to base 25 (for mapping reasons that I will explain later):

{7 13 15 2 3 13 22 9 4 16 0 16 24 13 2 22...}

To this vector, we will apply a transformation that will make it self-similar. For this process, I use the Microsoft Visual C++ programming environment, employing the following recursive function:

```
/*+++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
  Fractalization of an integer sequence:

  a(2n) = a(n); a(2n+1) = nth value of the sequence;
+++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++*/
```

This operation results in the following vector V, which I will use for the mapping process:

V: {7 7 13 7 15 13 2 7 3 15 13 13 22 2 9 7 4 3 16 15 0 13 16 13 24 22 13 2 2 9 22...}

The choice of mapping (using the modulo operation in this case) to two sets of musical parameters (pitch and note duration) is:

- **N: Pitch (Sound Height):** A set of 25 elements ranging from C#3 (0) to C#5 (24), chromatically ascending. Hence the choice of the base 25 mentioned earlier.
- **R: Note Duration:** A set of 10 elements (0-9), where 0 corresponds to a sixteenth note, 1 to an eighth note, and so on up to 9, which would be a half note tied to an eighth note.

Mapping the vector V to the sets N and R and transcribing them into musical notation results in this motif:

If we analyse it, we will see that it has a self-similar structure (repetition at different scales) and is therefore a musical fractal:

If we take the even-numbered notes, we see that they match the initial melody, but at a slower tempo. If we repeat the operation (now choosing the even-numbered notes from this new melody), we once again obtain the initial melody, even slower:

And so on. Naturally, the combination of different vectors and sets of musical elements will yield different motifs, thus creating an infinity of melodies waiting to be discovered. But we must not forget that, until they are inscribed into the timeline, they are merely timeless structures. Almost all of my works contain self-similar motifs; here are some examples:

"Melussyne" (Metaludios for piano, Book V).

"Nabla," for violin, clarinet, horn, piano, and electronics.

"Chigaday" for symphony orchestra.

MUSICAL RAW MATERIAL FROM FRACTAL IMAGES

Another technique that is frequently present in my works is the conversion of images (often of fractal nature) into musical material. This process determines not only the pitch of the notes, rhythmic patterns, and dynamics but also the overall form and narrative of the fragment. The musical raw material and timeless structures generated can be quite complex and extensive, involving the use of various computer programmes we've mentioned before.

The first step is the selection of an image that satisfies us. To do this, I use the freeware program Apophysis, which generates "flame" fractal images. This is a personal choice, and in principle, any fractal image generation program (like Ultrafractal) can be equally valid. The process of choosing the image itself is a trial-and-error-based search. Let's take, for example, the image I used for one of my Metaludios for piano, "Izar iluna.":

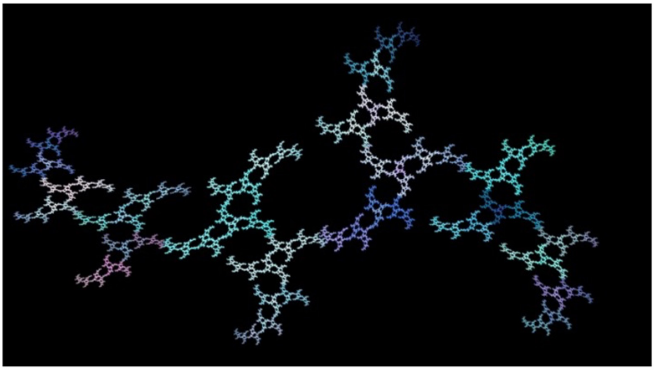

The next step is to convert the image into sound using the program *AudioPaint*:

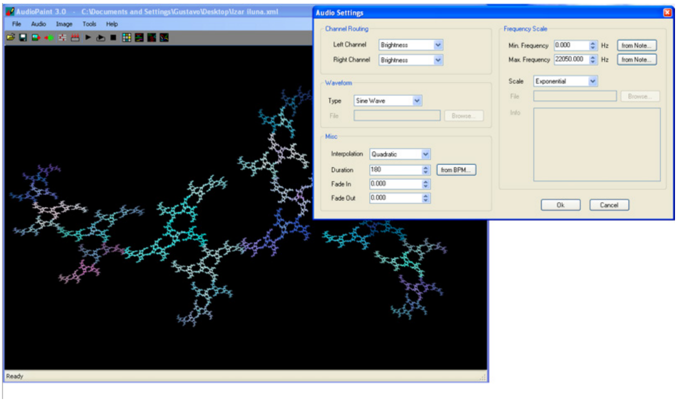

The X-axis in the image is mapped to the total time we desire, in this case, I chose 3 minutes.

The Y-axis in the image is mapped to a frequency range (in Hz). For this particular work, I selected a range of 100-1000 Hz. Of course, each choice determines the final outcome, which is why it's a trial-and-error process, involving listening, discarding what doesn't satisfy us, and experimenting with different parameters.

The color of each pixel determines the dynamics, with black [0,0,0] representing silence and white [255,255,255] representing maximum dynamic amplitude. The program allows us to convert the image into sound using pure sinusoids or granular synthesis. It offers many other options that can't be discussed in depth here. The process functions like massive additive synthesis, where each pixel of the image is transformed into an oscillator (or grain) with a frequency, duration, and dynamics based on the chosen parameters.

The result can be saved as a WAV file. A spectral analysis of the sound generated from the chosen image provides the following result:

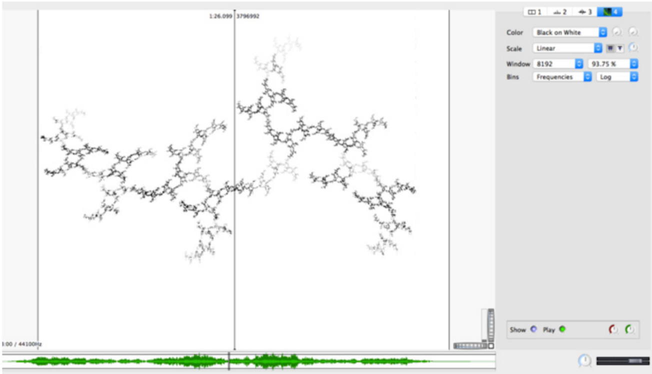

Which, logically, corresponds to the original image.

The next step is to convert the audio to MIDI so that I can work in the notation software. For this, I use Melodyne. This is the *Melodyne* work session corresponding to the audio generated by *AudioPaint*:

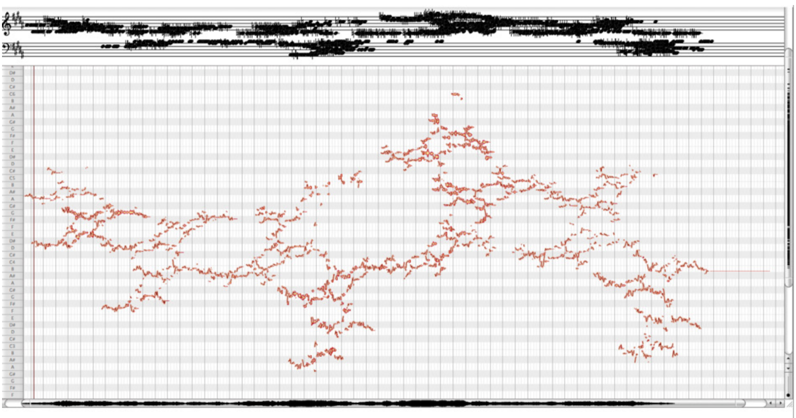

As we can see, the contour is identical to the initial image. Since it's a piano piece, we need to adjust the pitch to the nearest semitone (12ET):

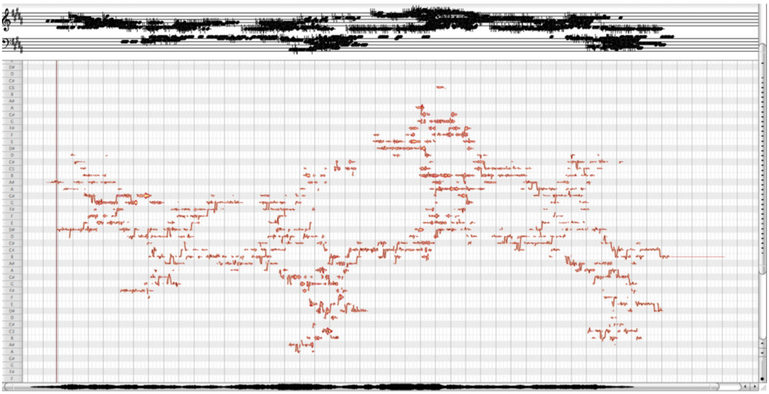

The contour remains recognisable, although it is slightly "flattened" when this operation is applied. Next, we generate a MIDI file that serves as the raw musical material to compose the piece:

This is a snippet of the raw final result, the timeless structure that we will work on to transform it into a work of art. From this point on, what I consider the most creative and interesting phase comes into play. It involves shaping the material, embellishing it, decorating it, adapting it to the instrumental possibilities, all while retaining its original essence:

An spectral image of the live performance audio of the final piece demonstrates that the fractal structure from the initial image is present:

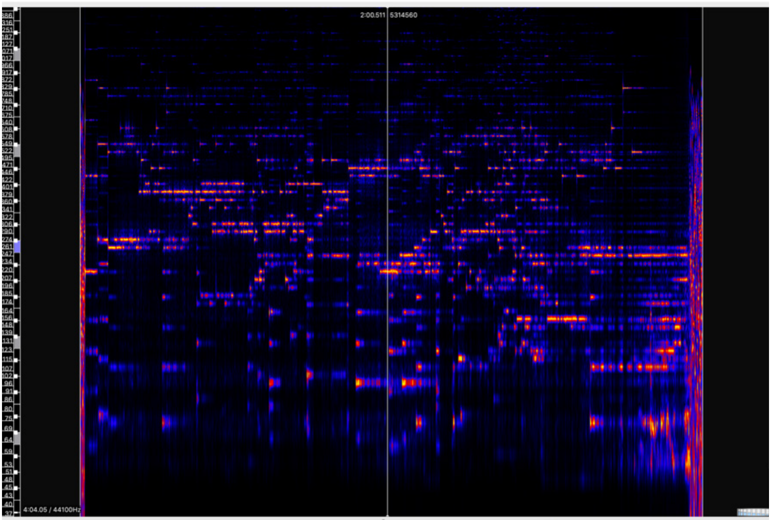

In the case of Izar iluna, the entire piece is generated from the original image. In other works, this technique is used for only a section within the temporal and narrative framework of the piece.

Here are other works in which I have used this technique:

Critical Strip, for string quartet (excerpt)

"Erbane" for Orchestra (Fragment)

“MUSICAL SLOW MOTION” AND “OVERPAINTING”

This compositional technique involves starting with a relatively short sound and, through specialised software, stretching it while maintaining all its richness (timbral complexity, dynamic range, etc.). For this purpose, I use the program Paul’s Extreme Sound Stretch, which I mentioned earlier. It’s a wonderful and user-friendly tool that provides us with a great deal of creative freedom thanks to the multitude of parameters we can control. The resulting sound can be used as a purely electronic track or, as we will see, to extract musical raw material that we will later shape with acoustic instruments.

In my Metaludio for piano, “Cassini’s Dream,” I used a 15-second recording published by NASA of radio emissions from Saturn’s magnetic field captured by the Cassini-Huygens probe. Here is the spectrum of this sound:

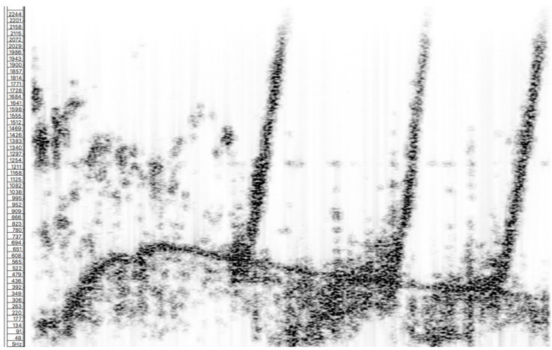

The stretched sound lasts 5 minutes and 45 seconds (345 seconds) and is therefore 23 times slower. Here is the spectrum of this new sound:

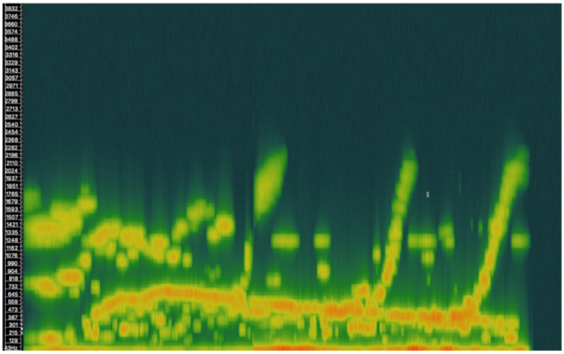

The stretching process is itself a creative and sonic exploration, given the immense flexibility of the software. I later load this stretched sound into *Melodyne* to extract the musical information, setting a tempo of quarter=60:

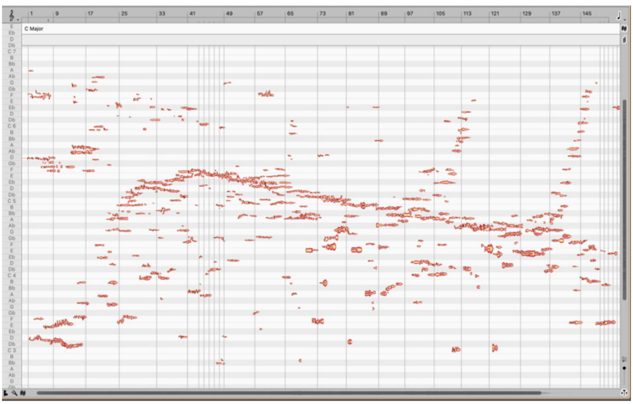

As before in Izar iluna, the pitch of the sound must be adjusted to the nearest semitone. Finally, a MIDI file is generated, which will be the raw material with which I will shape the composition:

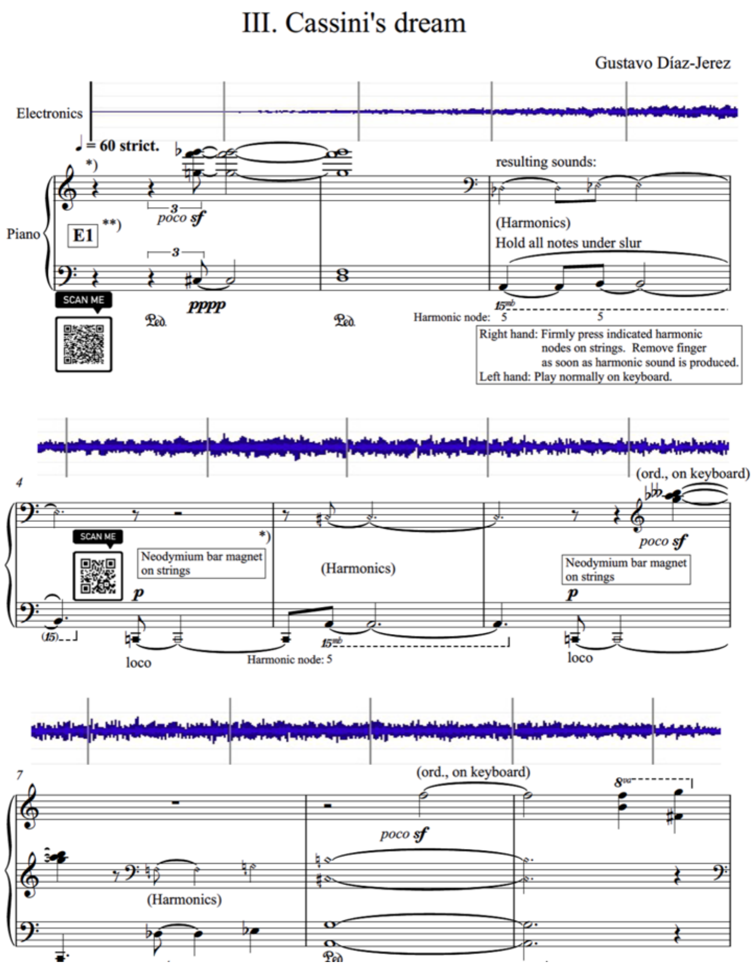

For this piece, the musical material has been crafted using numerous extended techniques (harmonics, clusters in the strings, the use of a neodymium magnet, etc.). The piece is performed simultaneously with the elongated electronic sound track. In this way, the live piano "paints" over the electronic sound since the musical material is directly derived from it. Because the tunings are not exactly the same (the electronic part is pure frequency without tempering), a sonically rich amalgamation is created. Additionally, the synchronisation of the electronic and live piano is never exact, making each performance slightly different. The result is a sonic analogy to the "overpainting" technique of the German painter Gerhard Richter, who, during a period of his career, used photographs as a base to paint over, embodying the interface between the representation of photographic content and abstract painting.

Another work in which I've employed this technique is "Tombeau de Perseus." In this case, the original sound is a 30-second recording from NASA originating from the supermassive black hole at the center of the Perseus galaxy cluster. "River of Oblivion" (Metaludios for piano, Book VI) combines this technique with image-to-sound conversion. The elongated sound comes from the conversion of an image using *AudioPaint*, as we saw earlier. In this case, it wasn't a fractal image but rather a photograph by the Venezuelan artist Blas González, who resides in Vigo.

Tombeau de Perseus, for Ensemble with Electronics.

River of oblivion (Metaludios for piano, Book VI)

“FRACTALIZATION” OF EXISTING MUSICAL MATERIAL

For this compositional technique, I retrieve existing material, usually from works of other composers whom I admire. This process is done entirely in the score editor, in my case, *Sibelius*. The fragment we are going to fractalise to illustrate the process belongs to a piano work by the Canarian composer Teobaldo Power. It’s the “arrorró,” a lullaby from the Cantos Canarios.

From now on, this fragment becomes raw musical material, a timeless structure that we will use to create a new work. The technique may appear simple, but it is capable of generating great complexity. It involves stacking several “layers” of the original material with affine transformations. In geometry, an affine transformation consists of a linear transformation followed by a translation. In general, an affine transformation is formed by linear transformations (rotations, homotheties, and shears) combined with a translation or displacement. In music, these operations are equivalent to rhythmic augmentation or diminution, transposition, retrogradation, inversion, retrograde inversion, and so on. Affine transformations are commonly used to construct fractals:

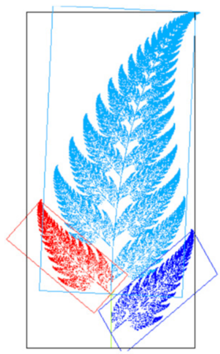

The following example uses three layers of Power’s original material: the first layer is the original material played 8 times slower and transposed to B-flat minor, occupying the entire duration of the piece:

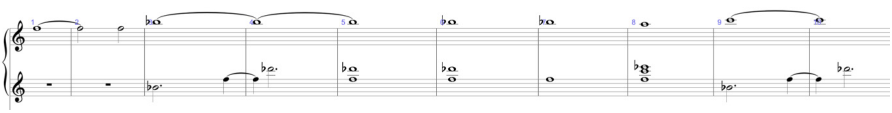

The second layer is the original material played 4 times slower, transposed to F minor, and shifted one measure plus an eighth note compared to the first layer:

The third and final layer is a transposition plus inversion of the original material, 2 times slower, shifted by 11 bars plus a sixteenth note relative to the first layer:

The three layers together would look like this:

We proceed to “sum” these three layers vertically, giving rise to the raw musical material:

Once this material has been worked on in a “traditional” way, that is, using the instrument, adding dynamics, phrasing, pedal indications, agogics, etc., we arrive at the final result, the work of art:

The possible variations allowed by this technique are virtually limitless. For "Orahan," I decided to use three layers of the original material, but in other works, I've used 6 layers (as in "Hommage an J. Brahms," from Metaludios, Book III) or even 12 layers (as in "Ommagio a Carlo Gesualdo," from Metaludios, Book IV). The type of affine transformation chosen for each layer and its placement in the timeline profoundly determine the final result. These decisions, as I've mentioned throughout this article, involve a process of trial and error, searching, listening, and discarding. Experimenting with different combinations and seeing what sonic result they produce (a process that is computationally irreducible in itself) is one of the most exciting experiences in the creative process.

Eine Hommage an J. Brahms (Metaludios, Book III)

Omaggio a Carlo Gesualdo (Metaludios, Book IV)

Pavana triste (tribute to Antonio José) (Metaludios, Book V)

Artificial intelligence (AI) is causing a revolution in all aspects of our lives. The progress of AI has been spectacular in the last 5 years, and artistic activities are not exempt from this revolution. AI can be an invaluable tool in the creative process.

During the COVID-19 lockdown between March and June 2020, I decided to experiment with *Magenta*, an open-source research project by *Google* that explores the role of AI as a tool in the creative process.

Machine learning algorithms learn from a dataset provided by the user. In my case, this dataset was the complete corpus of Metaludios for piano that I had up to that point, which consisted of 28 pieces. Once the system was configured, a task that was not easy and had to be done through the command line, the training began, lasting approximately 48 uninterrupted hours.

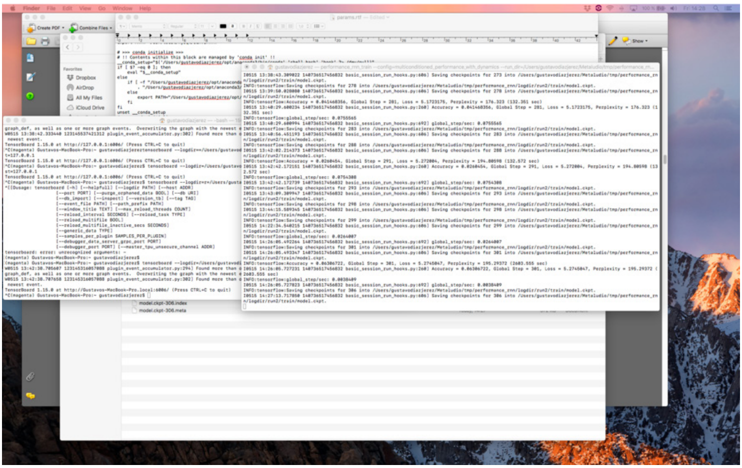

Training space

We must consider that the effectiveness of this process is proportional to the available computational power (memory, CPU, GPU, etc.), which, in my case, was very limited as I only had my laptop. Despite this limitation, the result was truly surprising, as we will see later on. During the training process, certain parameters need to be monitored, especially to prevent overfitting. When the model fits too closely to the dataset, it won't generalise well with new data. If this happens, it won't be able to perform the prediction tasks for which it was designed. In other words, we want the model to learn from the dataset and be able to generate new musical material based on it, imitating but not copying it literally. In real life, it's like teaching a composition student to write music in my pianistic style, but without directly copying anything from the samples they've worked with.

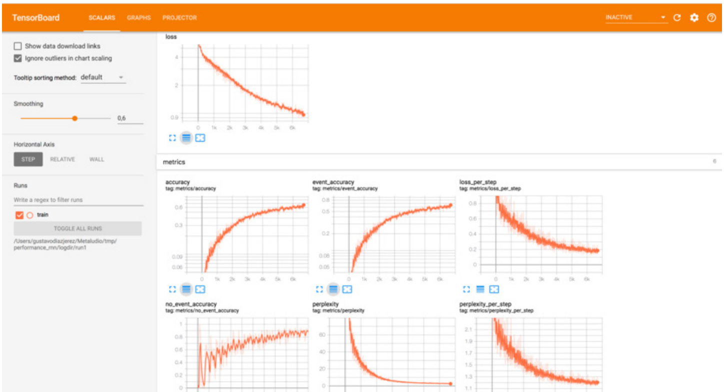

Training curves

Once the training process is completed, the most interesting part begins, which is starting to generate material from the resulting model. The model generates MIDI files of the specified duration. It's true that the MIDI format doesn't allow for all the expressiveness of musical notation that we might desire, but the system doesn't provide another option. In any case, it contains enough information (tempo, note pitch, rhythmic values, dynamics (MIDI velocity), sustain pedal information, etc.) for the generated segments to have sufficient musical richness. I generated around 300 MIDI files ranging from 4 to 5 minutes in duration. The next step was a process of selection, evaluation, as if it were a composition competition. The decision to accept or discard a segment was determined solely by listening (aesthetic judgment) and

difficulty (pianistic feasibility). From these initial 300 segments, I selected about 30 (promising ones). After a second listen, I kept around 14 (candidates), repeating the process. From these 14, I chose 7 (finalists), and from these 7, 3 (the chosen ones):

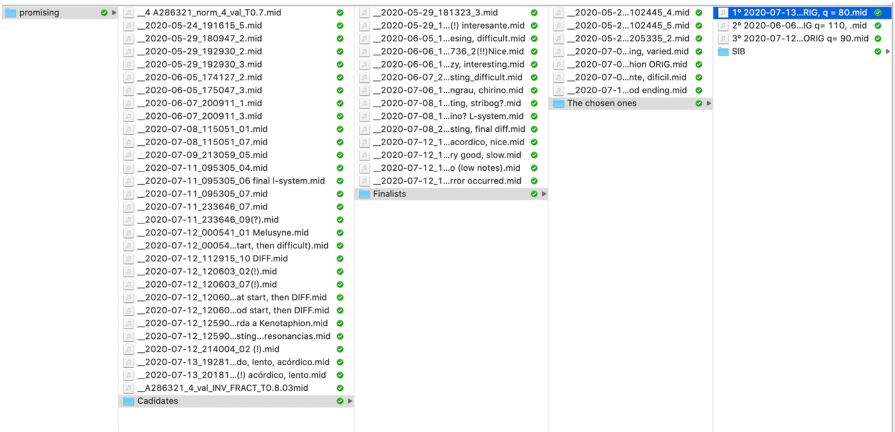

Selection

I initially thought, I thought about using this material as I usually do, as structures outside of time that I would then freely mould. However, in the end, I decided to leave them as they were, without modifying or changing anything, as a proof of concept. Nevertheless, I needed to rewrite the original MIDI files in *Sibelius* to make the scores more comfortable and clear to read. This involved a huge effort of listening and analysing the MIDI files to transcribe the dynamics and pedal indications into musical notation. I also added some phrasing and expression indications based on the emotions that the music suggested to me, but without changing anything that the machine had generated. Here is a fragment of the final scores, whose titles are related to concepts from AI.:

II. Long Short-Term Memory

Gustavo Díaz-Jerez

IV. Forget gate

Gustavo Díaz-Jerez

V. Hidden states

These pieces are now part of my Metaludios for piano. I must admit that I see myself in them. I have found multiple reminiscences of other Metaludios, although never explicitly (literally). In a way, the model is a kind of "apprentice" that has thoroughly studied the master's style and is capable of "creating" works in that style. Many profound philosophical questions and uncertainties arise here. Whose authorship do these works belong to, the algorithm or the one who trains the algorithm? After all, the dataset was provided by me, I configured and trained the model, and ultimately, I decided which fragments passed the aesthetic filter and became part of my catalog. As AI continues to advance and computational resources grow, the boundaries between

human creativity and silicon-based creativity will inevitably blur. But I firmly believe that, when used responsibly, AI will be a tool that opens up new paths and helps us be more creative, exploring new musical territories.

An artist must remain faithful to their own nature, which, from my point of view, is that attraction to something they've been drawn to since childhood. This starting point develops through learning, both intellectually and intuitively, leading to the creation of a personal method, a musical language that identifies us. The search for musical material in the computational universe has no limit. I always discover new structures and processes that yield immense compositional variety. And that, precisely, is a source of inspiration. I feel that this triangle of "search-model-result" is different every time, new in each work, infinite, like the scientific and the artistic world.